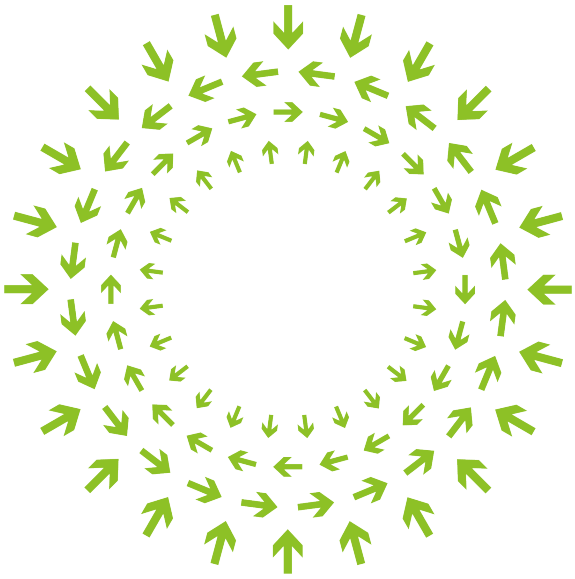

OPEN

Compute Project

# Hardware Management
# ICAP DRAM
Version 0.03
# Draft

Author: Hank Bruning, hank@jblade.com


# 1 Scope

This document defines the technical specifications identifying DRAM slots and modules used in Open Compute Project servers, storage devices and network switches. The specification is limited to the data format and commands defined in Intelligent Platform Management Interface specification and does not require the presence of an operating system on the device that is managed.

# Table of Contents

1 Scope....2
2 Overview....5
2.1 License....5
2.2 Reference Documents....5
2.3 Keywords....6
2.4 Out of Scope....6
2.5 Private Enterprise Number....6
2.6 OCP Specification Identification....6
2.7 FRU Records....7
3 DRAM Management....8
3.1 Inventory....8
3.2 Monitoring....9
4 Memory Module Slots....10
4.1 Logical Slot Identification....10
4.2 Slot Sensors....12
5 Memory Modules....14
5.1 Module Identification....14
5.2 Temperature Sensor....15
6 ID Assignment....17

# Revision History

| Date | Revision | Description |
|---|---|---|
| February 11, 2014 | 0.02 | Initial revision. Added content from document “Type A IPM Controllers” |
| June 9, 2014 | 0.03 | Added Section OCP Specification Identification. |

# 2 Overview

This describes the Intelligent Platform Management Interface (IPMI) capabilities(ICAP) implemented by an IPM Controller. It extends the IPMI 2.0 specification allowing Data Centers System Managers to implements a uniform identification and monitoring of DDR3 and DDR4 modules used in servers, storage devices and network switches.

This specification is one in a series of IPM Controller Capabilities(ICAP) which add functionality not found in the IPMI 2.0 specification.

This specification does not contain any requirements for hardware dimensions, connectors or electrical interfaces.

## 2.1 License



## 2.2 Reference Documents

These documents are referenced by this specification.

### 2.2.1 Specification Documents

| Acronym | Date | Specification |
|---|---|---|
| DDR3 SPD | 9/1/2011 | JEDEC Standard No. 21-C, Annex K: Serial Presence Detect (SPD) for DDR3 SDRAM Modules, Release 4 |
| DDR4 SPD | November 2013 | JEDEC Standard No. 21-C Release 23, Annex L: Serial Presence Detect (SPD) for DDR4 SDRAM Modules |

| Acronym | Date | Specification |
|---|---|---|
| DDR Thermal Sensor | November 2009 | JEDEC Standard No. 21-C Release 19, Section 4.1.5 TS3000 Standalone Thermal Sensor Component |
| IPMI 2.0 | 10/1/2013 | Intelligent Platform Management Interface Specification Second Generation v2.0, Document Revision 1.1 |

## 2.3 Keywords

**shall**

A keyword indicating a mandatory requirement; designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this specification.

**shall not**

A keyword used to describe a feature, function, or coded value that is defined in a specification to which this specification makes a normative reference where the use of said feature, function, or coded value is not allowed for implementations of this specification.

## 2.4 Out of Scope

This specification does not contain any requirements for hardware dimensions or connectors .

Ethernet connection or the TCP/UDP protocol is not required to implement this specification.

The electrical interface between the IPM Controller and any device is not defined in this document.

Hot plug DDR memory modules are not supported.

## 2.5 Private Enterprise Number

The IPMI Commands and FRU records defined in this document utilize the Private Enterprise Number 42623 assigned to OCP by the Internet Assigned Number Authority , www.iana.org. In a twist of fate that only a IPMI implementer will enjoy, the number assigned to OCP ends in the IPMI UDP port number, 623.

## 2.6 OCP Specification Identification

This specification is identified by five fields defined in the OCP document Hardware Management SPEC ID. The five field are present in Table 2, OCP Specification Descriptor byte offsets zero to five and Table 4, Get OCP Specification Version byte offsets one to six. The value of the five fields are found in the following requirements.

IPMA-IPM-2.1 The *OCP Specification ID* used to identify any revision of this document **shall** be a value of 0x2.

IPMA-IPM-2.2 The *OCP Specification Revision* used to identify this specific document **shall** be a value of 0x0.

## 93 2.7 FRU Records

94 All FRU records defined in this specification contain the Private Enterprise Number as
95 the first three bytes after the record header checksum and is written Least Significant
96 byte first.

# 3 DRAM Management

This specification is part of the IPMI Capabilities(ICAP) Identification and Status(IDS) series providing a uniform interface for managing servers, storage devices or network switches. ICAP enables a Data Center System Manager to control rack mounted hardware with minimum operator intervention. The ICAP specifications allow a single technician to operate 10,000 to 50,000 servers.

The IDS set of specifications allow the physical location of a device in a two dimensional grid and provide status information which includes sensors for device present/absent, temperature and voltage.

This document, identified by acronym ICAP-DRAM, defines set of requirements for an IPM Controller implementing the Intelligent Platform Management Interface (IPMI 2.0) to monitor DDR3 and DDR4 modules on Servers, Storage Devices and Network Switches, which this document collectively calls IT hardware. No distinction is made between the different types of IT hardware. The benefits to the Data Center operator are a uniform, vendor and hardware neutral methods for:

- inventory data collection of DDR3/DDR4 manufacturer id, model, serial numbers
- on-site customer acceptance tests validating server memory size
- System Manager error messages identifying the specific DRAM slot
- identification of empty DRAM slots so servers can be upgraded
- identification of DDR3/DDR4 memory capacity and module errors
- report temperature alarms on DDR3/DDR4 modules

The requirements in this document are specific to IPM Controllers in IT hardware and no requirements are made on the System Manager itself. If a System Manager is not present, the IT Hardware will function normally.

This specification is not mandatory for IT Hardware. There is no architectural limitation preventing a single chassis from containing IPM Controllers implementing this specification and IPM Controllers implementing proprietary IPM Controllers.

DDR3 and DDR4 memory modules which may be removable or be soldered to a PCB. The FRU Information Records have been designed so DDR3/DDR4 memory modules that are soldered to a board can be coexist with removable memory that on the same circuit board. This allows a System Manager to determine the correct amount of memory when a single CPU uses low cost soldered DDR memory modules and the higher cost removable modules.

## 3.1 Inventory

An IPM Controller allows IT hardware a common interface to identify the location and inventory of memory modules supported by the main CPUs. A System Manager query the IPM Controller to determine what type of memory modules, if any, are installed in the memory module slots. This allows a System Manager to maintain a real time database of the type, memory capacity of each memory module and the total installed memory capacity available to each CPU.

This allows a System Manager to:

1. detect failed or incorrectly seated memory modules. The System Manager's database contains the expected number of modules in server and can detect if the actual number is lower.
2. detect which servers can be upgraded to higher memory capacity by populating unused slots or replacing low capacity modules with higher capacity modules.
3. compare memory module thermal design parameters to choose lowest cost cooling solutions. For example a Data Center may find it more cost effective to replace 85° C maximum operating temperate memory modules with 95° C modules than to redesign air conditioning systems.
4. select hardware based on actual memory capacity to install operating systems and the applications.
5. predict the number of memory modules which can be reallocated to other IT hardware when permanently deactivating running systems. For example when RAID is going to be removed from the Data Center how many DDR3 modules can be reallocated to compute servers.

## 3.2 Monitoring

One goal of the ICAP specifications is to provide Data Center operators information on how close the IT hardware is the operating limits. This information is valuable when raising air temperature in a data hall in an effort to decrease operating costs.

DDR3 memory modules rated at 95° C are required to have temperature sensors. The sensors are optional for 85° C modules.

A System Manger communicating with a IPM Controller implementing ICAP-DRAM(this specification) can determine, when a memory module is withing the vendor defined temperature envelope.

# 4 Memory Module Slots

A location where a memory module may be present is called a slot. It may be a card edge connector or a slot is memory soldered to a printed circuit board.

This specification is not designed to allow memory module slots or memory modules to be inserted or removed while an IPM Controller is operating.

## 4.1 Logical Slot Identification

An IPM Controller identifies the memory module slots associated with the main CPU(s) with an IPMI FRU Device Locator Record. It identifies the FRU address that represents the memory module slot. The *Device ID* Field contains the ASCII characters that uniquely identify the slot either in manufacturer documentation or the characters are printed on the circuit board adjacent to the memory slot. This allows a System Manager to create error messages that are specific to a memory module slot.

Each LUN ID has a separate Device SDR repository. A single Device SDR Repository contains all bays FRU Device Locator Records that are managed by the LUN ID. There are a maximum of four LUN IDs within an IPM Controller.

**Illustration 1: IPMC to Device SDR Repository Mapping**

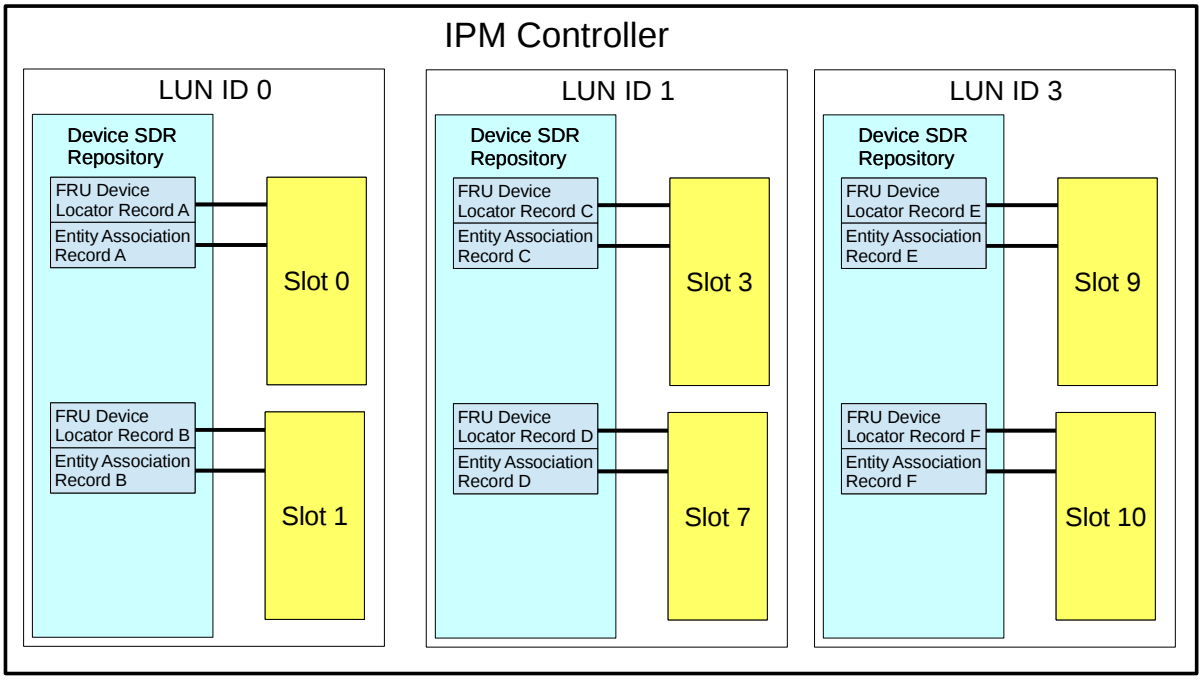


Each memory module slot's FRU Device Locator record contains an Entity ID Instance

Number that is unique from all other memory module slots' FRU Device Locator records within the same Device SDR Repository.

DRAM-SLO-4.1 Each DDR3 or DDR4 memory module slot shall have IPMI FRU Device Locator Record present in the Device SDR Repository present in the same LUN ID as found in the FRU Device Locator Record *Logical-Physical / Access LUN / Bus ID* Field(byte8) bits 4:3.

DRAM-SLO-4.2 Each IPMI FRU Device Locator Record for a memory module slot **shall** have the Fru Enity ID field(byte 13) set to 8h.

DRAM-SLO-4.3 If the memory slot is for a DDR3 module the FRU Device Locator Record for a memory module slot **shall** have the *Device Type* field(byte 11) set to a value to C0h( Table 3: IPMB/I2C Device Type Codes).

DRAM-SLO-4.4 If the memory slot is for a DDR3 module the FRU Device Locator Record for a memory module slot **shall** have the *Device Type Modifier* field(byte 12) set to a value to C0h( Table 3: IPMB/I2C Device Type Codes).

DRAM-SLO-4.5 If the memory slot is for a DDR4 module the FRU Device Locator Record for a memory module slot **shall** have the *Device Type* field(byte 11) set to a value to C1h( Table 3: IPMB/I2C Device Type Codes).

DRAM-SLO-4.6 If the memory slot is for a DDR4 module the FRU Device Locator Record for a memory module slot **shall** have the *Device Type Modifier* field(byte 12) set to a value to C1h( Table 3: IPMB/I2C Device Type Codes).

DRAM-SLO-4.7 All FRU Device Locator Record for a memory module slot **shall** have an unique combination of IPMI Channel, IPMB address, LUN ID and FRU ID.

DRAM-SLO-4.8 No Device Locator Records for a memory module slot **shall** have an identical value in the *FRU Entity Instance* field.

DRAM-SLO-4.9 All Compact and Full Sensor Records for the memory module **shall** be present in the same LUN ID's Device SDR Repository containing the memory module slot's FRU Device Locator Record.

### 4.1.1 Physical Slot Identification

Memory slots are identified with two or three characters. The identification is called a Slot ID. Slot IDs may appear on the printed circuit board, documentation or in System Manager error messages.

The leading character of a Slot ID is an uppercase letter indicating the row number assigned to the memory slot. Rows are assigned sequentialy increasing alphabetic characters. The first row is always “A” and is closest to the air inlet.

The trailing characters are a numeric value from 0 to 29 indicating the column number assigned to the memory slot. Columns are assigned from left to right. Columns are sequentially assigned a number with no numeric gaps allowed.

There is always a memory Slot ID A0.

Illustration 2: Example Memory Slot ID Layout

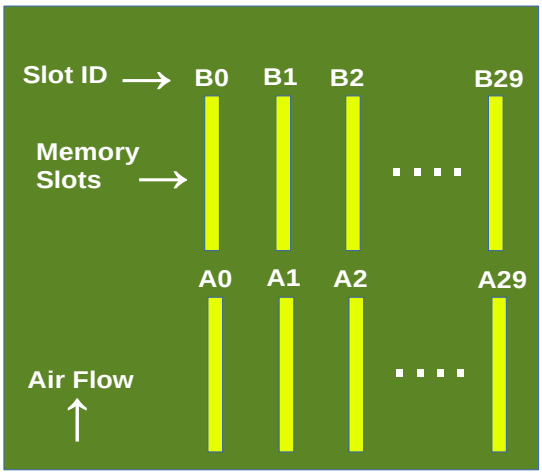


DRAM-SLO-4.10 An IPMI FRU Device Locator Record for a memory module slot **shall** have a unique *Device String* compared with all other memory module FRU Device Locator Records within the same Device SDR Repository.

DRAM-SLO-4.11 The first character in the FRU Device Locator Record Field *Device String* **shall** be an upper case letter indicating the row containing the memory slot.

DRAM-SLO-4.12 The second and third character in the FRU Device Locator Record Field *Device String* **shall** be a numeric indicating the column containing the memory slot.

DRAM-SLO-4.13 The maximum length of the Device Locator Record Field *Device String* **shall** be three ASICII characters.

DRAM-SLO-4.14 If the column number of the memory slot is less than ten the length of the Device Locator Record Field *Device String* **shall** be two ASICII characters.

DRAM-SLO-4.15 The IPM Controller **shall** identify memory slot column from left to right with a digit with the left most digit being zero.

DRAM-SLO-4.16 The IPM Controller **shall** identify the left most memory slot column as digit zero.

DRAM-SLO-4.17 The IPM Controller **shall** identify the memory slot row closest to the air intake with the upper case letter A.

## 4.2 Slot Sensors

Memory modules can not be inserted and removed while the IT hardware is running. The IPM Controller is responsible detecting optical modules when the IPM Controller is reset or powered on. It generates sensor SDR records and inserts them in the Device SDR repository.

OPTI-BAY-4.18 The IPM Controller **shall** detect the memory modules when it is powered on or reset.

The Entity Instance number present in the FRU Device Locator Record for a single memory module bay is identical for all Compact and Full Sensor Records related to a single memory module slot and it's memory module.

### 4.2.1 Module Presence Sensor

Each memory module slot has a IPMI sensor to detect when the memory module is present. The sensor is present in the Device SDR repository at all times.

DRAM-SLO-4.19 For each memory module slot FRU Device Locator Record there **shall** be a Presence Detect Compact Sensor Record with the *Entity Instance*(byte 10) field identical to the FRU Device Locator Record's *FRU Entity Instance* Field(byte 14).

DRAM-SLO-4.20 The memory module slot Presence Detect Compact Sensor Record and Memory Module Slot FRU Device Locator Record that have identical values for the Entity Instance **shall** be present in the same LUN ID's Device SDR.

DRAM-SLO-4.21 Each memory module slot Presence Detect Compact Sensor Record s**hall** have the *Entity ID* field(byte 9) set to 8h.

DRAM-SLO-4.22 Each memory module slot Presence Detect Compact Sensor Record s**hall** have the *Event/Reading Type Code* field set to 8h, Device Absent/Device Present.

DRAM-SLO-4.23 When an memory module is removed all Full and Compact Sensor records **shall** be removed from the Device SDR prior to the IPM controller setting the memory module slot's Presence Detect sensor state to 0h(Device Removed / Device Absent ).

DRAM-SLO-4.24 When an memory module is inserted all it's Full and Compact Sensor records **shall** be inserted into the Device SDR prior to the IPM controller setting the memory module slots Presence Detect sensor state to 1h(Device Inserted / Device Present).

DRAM-SLO-4.25 The Compact Sensor Record defining a memory module slot Presence Detect sensor **shall not** be removed from the Device SDR Repository.

# 5 Memory Modules

This specification has no requirements on the interface between the IPM Controller and the memory module. It may be proprietary. One method is an I²C interface between the IPM Controller and a memory module slot. The I²C interface is defined by the JDEC C-21 DDR Thermal Sensor specification. It allows up to eight DDR memory modules per I²C bus.

The I²C interface to each memory allows an IPM Controller to detect if a DDR3 or DDR4 memory module is present in the slot. Configuration information and operating status can be read from the memory module.

## 5.1 Module Identification

An IPM Controller maps each memory module's Serial Presence Detect(SPD) data bytes to a different FRU ID's FRU Information area as a OCP defined SDR Multi Record. An IPM Controller is not required to read the SPD data until it received an IPMI Read FRU Info command for the FRU ID. Once the IPM Controller or the System Manager has read the SPD data no changes are allowed to that data until the IPM Controller is rebooted.

The format of the SPD data is not compatible between the DDR3 (DDR3 SPD Appendix K) and DDR4 (DDR4 SPD Appendix L). They share a common format for Byte 2 of the SPD data that differentiates between the types of memory. Note that the DDR4, Appendix L is only present in the JEDEC Standard No. 21-C Release 23 and later versions.

The the SPD data contain a Cyclical Redundancy Code(CRC) . The IPM Controller is not required to validate that the CRC is correct.

### 5.1.1 DDR3 SPD

The DDR3 SPD specification documents a sequence 128 bytes on each DDR3 module that identify the size, capabilities, vendor and serial number.

The thermal characteristics of the memory module are defined by JDEC 21-C as:

- Normal temperature range, 0 to 85 degrees C.
- Extended temperature range, 0 to 95 degrees C.

This information is present in the SPD as field *SDRAM Thermal and Refresh Options*(byte 31).

### 5.1.2 DDR4 SPD

The DDR4 SPD specification documents the of sequence 512 bytes on each DDR4 module that identify the size, capabilities, vendor and serial number.

The thermal characteristics of the memory module are defined by the hardware vendor an present in the SPD as field *SDRAM Thermal and Refresh Options*(byte 8).

### 5.1.3 DRAM Multi Record Description

An IPM Controller maps each memory module's Serial Presence Detect data to a single FRU's Multi Record area containing a DRAM Module Description Record.

244 **Table 1. DRAM Module Description Record**

| Offset | Field Length | Field Name |
|---|---|---|
| 0 | 1 | *Record Type ID*. For all records defined in this specification a value of D0h (OEM) is used. |
| 1 | 1 | End of List/Version<br>[7:7]- End of List. Set to one for the last record.<br>[6:4]- Reserved. Write as 0h.<br>[3:0]- Record Format Version. For this specification 0h. |
| 2 | 1 | Record Length |
| 3 | 1 | *Record Checksum*. The zero Checksum of the record. |
| 4 | 1 | *Header Checksum.* The zero Checksum of the header. |
| 5 | 3 | *Manufacturer ID*. The Private Enterprise number assigned to OCP. Write as 42623 (A67Fh) . Least significant byte first. |
| 8 | 1 | *OCP Record ID*. 04h |
| 9 | 1 | *Record Format Version*.<br>[7:4] Reserved. Write as 0h.<br>[3:0] Format Version ID. Use 0h to identify this table. |
| 10 | 9+N | *DRAM Module Description*. The memory module Serial Presence Detect data formatted according to DDR3 SPD or DDR4 SPD specification. The size, N, is determined from the SPD data byte offset zero, bits [3:0]. |

DRAM-MOD-5.1 If a memory module is detected in a memory module slot the IPM Controller **shall** set the FRU ID's Common Header field *MultiRecord Area Starting Offset* to the starting offset of the FRU Multi record containing the DRAM Module Description

## 245 5.2 Temperature Sensor

246 The DDR3 and DDR4 SPD data contains a field indicating the presence of a
247 temperature sensor on the memory module. The sensor is optional for 85° C
248 modules and mandatory for 95° C modules. When the sensor is present this section
249 documents how an IPM Controller interacts with the sensor. The sensor allows each
250 DDR memory module to provide temperature readings that provide information on
251 the health and status of the hardware. If a server is implemented without Fan speed
252 sensors the rise in DDR temperatures is often the first indication that Fan is failing.

253 The DDR Thermal Sensor specification defines the sensor thresholds and how the
254 thermal sensor is read.

255 When the memory module is detected in a slot the IPM Controller reads the memory
256 modules SPD data to determine if the module contains a temperature sensor. If a
257 temperature sensor is present the IPM Controller copies the Entity Instance Number
258 from the FRU Device Locator record and creates temperature sensor thresholds to
259 assemble the temperature sensors Full Sensor Record. Then the Full sensor record is
260 inserted into the same LUN ID's Device SDR repository where the FRU Device Locator
261 Record was found.

DRAM-MOD-5.2 If the IPM Controller detects that a memory module implements a temperature sensor the IPMI Controller **shall** create a Full Sensor Record with the *Entity ID* field set to 8h(memory module) in the Device SDR Repository of the LUN ID field(byte 7) found in the Full Sensor Record.

DRAM-MOD-5.3 The memory module temperature Full Sensor SDR **shall** have the *Sensor Type* field set to 1h(Temperature),

DRAM-MOD-5.4 The DDR memory module temperature Full Sensor SDR **shall** shall have the upper non-critical threshold from the data in DDR Thermal Sensor, Section 6.5.1 Alarm Temperature Upper Boundary Register.

DRAM-MOD-5.5 The DDR memory module temperature Full Sensor SDR **shall** have the lower non-critical threshold from the data in DDR Thermal Sensor, Section 6.5.2 Alarm Temperature Lower Boundary Register.

DRAM-MOD-5.6 The DDR memory module temperature Full Sensor SDR **shall** have the upper critical threshold from the data in DDR Thermal Sensor, Section 6.5.3 Critical Temperature Register.

DRAM-MOD-5.7 The memory module temperature Full Sensor SDR **shall** have the *Entity Instance* field(byte 10) set to a value found in a memory module slot's FRU Device Locator Record's *FRU Entity Instance* Field.

# 262 6 ID Assignment

263 **Table 2: FRU Information Record ID Assignments**

| FRU Record Name | Table Number | Record ID | Chassis Manager | Node Manager |
|---|---|---|---|---|
| DRAM Module Description Record | 1 | 4h | Yes | Yes |

264 The OEM values for the IPM 2.0, Table 43-1, IPMB/I2C Device Type Codes.

265 **Table 3: IPMB/I2C Device Type Codes**

| Sensor Type Code | Sensor Name |
|---|---|
| C0h | DDR3 Memory Slot |
| C1h | DDR4 Memory Slot |